

The Logic of Google Ads

When should you buy ads? Let's assume your goal is for people to click on the ads and give you money. (Reasons this may not be true: persuasion, brand-building, budget-maximizing.) The return from a block of ads is thus revenue - marginal_costs - ad_costs. Ads are an investment like any other; you keep buying them until your return on investment (revenue - marginal_costs / ad_costs) equals your cost of capital (usually the interest rate).

For simplicity, we'll assume your marginal cost is zero. (My marginal cost is almost always zero, so this doesn't strike me as too unrealistic.) So how do you estimate revenue? You can track how much money people who click on your ad give you, but this has two flaws. First, customers often give you more money over time. Maybe they buy level one of your video game when they click on the ad, but then they may buy levels two and three the next day after they beat level one. The future is always in the future, so revenue-per-user numbers may be too small.

Second, they might have given you money anyway. Your video game ads probably run on video game review sites, where readers might buy your game just from the review, even if you hadn't bought an ad. So your revenue numbers may be too big.

But these problems aren't so serious. In the first case, the worst that happens is you don't buy as many ads as you should. In the second, you don't actually lose money, it's just that some extra profit you could have kept has gone into ads.

Let's turn to the ad seller. They probably want to maximize how much they charge per ad impression (CPM). (Reasons this may not be true: unseemly ads.) A good way to do this is to hold an auction. It's impractical to have everyone bid live, so Google auctions work like eBay auctions: you enter the maximum you're willing to pay and get charged just enough to beat the other bidders. (One can think of this as a computer-simulated auction where everyone keeps bidding up the price by pennies until they hit the maximum they're willing to spend.)

But what are you bidding on? Ad sellers want to maximize revenue per impression, but ad buyers want to maximize profit per expense. In an ideal world, ad sellers auction off impressions (this is [what Google Ad Manager does](#)) while ad buyers bid per dollar of profit (entering their cost of capital).

Determining how much profit you make from an ad is hard. Can we just trust you? Let's say you make \$2 in profit per 1000 impressions and everyone else makes \$1. Now you can lie and say you make \$1 in profit and then pay twice as much per profit-dollar. Now you pay the same amount as before, but you win all the profit-dollar auctions. Now that's not wrong — you're clearly making more money than the other bidders, so you should win — but your bid isn't cost-per-profit anymore, it's cost-per-impression.

What if you paid based on revenue? Verifying revenue is difficult, but Google could do it if everyone was using Google Checkout. (If you sent some of your users to a non-Google Checkout system, Google could catch you and fine you.) Google offers [nicer ads](#) to Checkout users, but they still don't have much market share, making this system impractical at present.

Some search engines apparently had cost-per-action (CPA) auctions, where you paid based on how many people actually bought things. I have no idea how they made that work, since lying about how many people took an action seems really profitable and easy. Maybe that's why no one does this anymore.

That just leaves cost-per-click (CPC). Cost-per-click seems ideal, since it's verifiable by both the ad seller (who uses a redirect link to track clicks) and the ad buyer (who sees the users show up on their page). It's a nice half-way point between buyer and seller.

So the ad seller holds an auction for CPC and multiplies CPC by click-thru-rate (CTR) to calculate CPM. They show the highest CPM ads, charging each the bidder below them's CPC, times their relative CTRs. (In reality, Google doesn't just use CTR; they also factor in the relevance of the ad and the quality of the page it goes to.) And, voila: we've derived the basics of an online ad system.

This works out great for the ad seller — they maximize CPM, just like they wanted — but the ad buyer is still stuck converting their ROI into CPC. The ad buyer, recall, wants to increase their spending on ads (now determined to be CPC) until their return on investment equals their cost of capital.

It seems like this should be pretty easy, and indeed Google does provide [tools to calculate ROI](#), but apparently not to optimize it. What they do provide is [a tool](#) to optimize your cost-per-action. **Does anyone know why this is?**

It seems like an automatic ROI optimizer would lead many people to spend more money on ads. It's hard to believe Google is leaving all that money on the table.

But Google does intelligently optimize the ads themselves. The variance in click-thru rates between different ads is huge — it's not uncommon to see two very similar ads, but one gets ten times as many clicks as the others. Google lets you put in as many ads as you like and [automatically rotates them](#), showing ads with better CTRs more often.

So far we've just had a single ad seller. In the real world, lots of people want to sell ads and lots of people want to buy them. How do you match them up?

One option is make the buyer choose. This is how Google Search works: Google holds an auction for each search query and buyers pick which ones they want to compete in. Another is to group related websites together and run ads evenly across all of them. This is how most smaller ad networks work. And then there's AdSense. AdSense scans a page for relevant keywords, then runs the Google Search ads that won auctions for those keywords.

Google also knows a lot about ad *viewers*. By tracking what web pages you visit, they know [what topics you're interested in](#). I'm apparently interested in Unix, the environment, elections, government, and social science, so Google prefers to show ads on those subjects to me.

But there's another way to think about ad matching: as a giant optimization problem. Which combinations of user, ad placement, and advertisement optimize click-thru rates (or, ultimately, ROI)?

For each of these, there are lots of variables. For each user, you know their history, geographical location, computer (browser, operating system, screen size), ISP, etc. For each ad placement, you know time of day, hosting website, page content, etc. And for each ad, there are numerous possible variations in phrasing and design that can be tested, as mentioned before.

The possible combinations are infinite. You can't test all of them, so you need to come up with ones that are plausible. You can look at which combinations worked in the past: has this ad done significantly better in some cities than others, or at some times than others? And you can look for patterns across ads: do ads that do well on CNN also do well on MSNBC? These hypotheses can then be tested and, if they work, you start running ads more there.

Netflix claims they've made millions from slight improvements in their movie recommendations.¹ When they offered a prize for more, researchers found thousands of tiny patterns and came up with all sorts of innovative algorithms to try to get an edge. After 32 months, researchers doubled the algorithm's effectiveness.

Imagine how much more is at stake for Google. Last year, they received \$21 billion in ad revenue, of which 60% was apparently profit. Even tiny improvements would be worth the highest salaries — a 0.004% improvement would make \$500,000. Doubling it would create unspeakable wealth.

Yet Google has no contest for improving ad click-thru rates. Indeed, press reports suggest they don't even have an internal team working on it. The AdWords user-interface (recently redesigned from jaw-droppingly wretched to just wretched) would seem to suggest they don't do this kind of optimization at all. Their blog asks people to [optimize things manually](#). No doubt there are some things humans (even ad purchase reps) can do better than computers, but surely there's a lot more they can do together — with humans giving the machine additional hints and hypotheses to test. But there doesn't seem to be anything like that.

It's hard to believe this is true. It's hard to believe this can last.

Google's chief economist [claims](#) that Google's sewn up the ad market by being better than everyone else. What if you made an ad network that was better than Google?

Right now Google takes a 20% cut of every auction price. What if you were willing to take just 10%? You could give ad sellers a slightly higher CPM — they'd gladly run your ads when they paid more and Google's the rest of the time. Then you can offer ad buyers a slightly lower CPC. As long as the money people made was more than the cost of setting things up, they'd switch. I'm actually not sure why this hasn't happened.

Now imagine that you were a genius CS student who could come up with a better ad optimization algorithm. Your system would have a higher overall CTR, since it presented users with better ads. This means that, again, you can pay higher CPMs (since more people click per impression). And you can redirect some of the money you would spend on higher CPMs into lower CPCs, to attract advertisers.

But to develop the algorithms and do the optimization you need the data. Lots of it — lots of users, lots of advertisers, lots of ad spots. No startup will ever have that; it's only left to Google (or whichever giant eventually replaces them).

I'm not normally one to be too concerned about improving Google's bottom line (they seem to be doing alright), but as an ad buyer I'm frustrated I have to do this work myself. I'd rather solve the problem for everyone. And if Google wants to pay me for that, I certainly wouldn't mind.

It's weird that Netflix is so much more interested in this than, say, Amazon. Amazon makes money on every sale, whereas Netflix loses money every time they send a DVD out. Netflix claims they make up for this in higher customer retention rates, but why didn't Amazon think of this first? ■

You should follow me on twitter [here](#).

November 3, 2009